Fault Modelling in System-of-Systems Contracts

Zoe Andrews, Jeremy Bryans, Richard Payne
School of Computing Science
Newcastle University, UK
{*firstname.lastname*}@ncl.ac.uk

Klaus Kristensen
Bang & Olufsen, Denmark
krt@bang-olufsen.dk

arXiv:1404.7775v4 [cs.SE] 7 Oct 2014

***Abstract*—The nature of Systems of Systems (SoSs), large complex systems composed of independent, geographically distributed and continuously evolving constituent systems, means that faults are unavoidable. Previous work on defining contractual specifications of the constituent systems within SoSs does not provide any explicit consideration for faults. In this paper we address that gap by extending an existing pattern for modelling contracts with fault modelling concepts. The proposed extensions are introduced with respect to an Audio Visual SoS case study from Bang & Olufsen, before discussing how they relate to previous work on modelling faults in SoSs.**

***Keywords* – systems of systems, modelling, architectural frameworks, contracts, faults.**

I. Introduction

Systems of Systems (SoSs) are an active area of research due to the complexity that is present as a consequence of several characteristics common to SoSs [1]; including: the independence of the constituent systems (CSs) making up an SoS, the distribution of CSs, the evolution and reconfiguration of an SoS, and the emergent behaviours of an SoS. In addition to these characteristics (and conceivably due to them), Kopetz states that faults – considered exceptional in traditional monolithic systems – are normal and to be expected in SoSs [2]. We therefore need methods to assist in understanding and designing for faults in SoSs.

The job of the SoS integrator is to integrate "disparate independent [constituent] systems up to an interoperable formation of systems" and it is often a "complex, risky, long, and frustrating effort" [3]. This is due in large part to "emergence": a natural feature of SoS [4]. We define *emergence* in accordance with Nielsen et al. as "the behaviours that arise as a result of the synergistic collaboration of constituent [systems]" [5]. Emergence may be desired or undesired: the SoS integration must produce the desired emergent behaviour while preventing undesired emergent behaviour within the SoS. Commonly CSs are not under the control of the SoS integrator, and thus we propose a contractual description of these CSs that captures their external behaviour [6], [7].

In this work, a contractual description (or *contract*) defines the behaviour of a CS in terms of its functionality and ordering of events. Currently, however, contracts contain no support for explicitly recording and dealing with faults. Here we extend the definition of contracts to enable the explicit recording of faults and fault tolerance in SoSs. We do this by building on two existing pieces of work: on contracts and on fault modelling.

We envisage these extensions being used by two groups: CS designers and SoS integrators. The designer of a CS (or its contract) may identify and acknowledge any deviations from desired behaviour of CSs. They may also state any such deviations from their environment (for example from other CSs in the SoS) that they mitigate. These extended contract definitions may be used by the SoS integrator as inputs to the integration process. The SoS integrator may also observe deviations of the CSs not under their control, and wish to mitigate them. Fault modelling techniques can be applied that explicitly state what deviations may be present in the CSs of the SoS, and identify those CSs that may be affected by or mitigate those deviations.

In this paper, we propose extensions to an existing method for defining contracts; providing traceable fault modelling techniques to understand the effects of faults in contracts. This allows: (i) CS designers to annotate their CS contracts with failure modes and fault masking capabilities; (ii) SoS integrators to describe faults identified in the CSs; and (iii) SoS integrators to define the means to recover from, or mitigate, the effects of SoS faults.

In Section II we go into more detail on relevant background on modelling contracts and faults before describing related work in Section III. In Section IV we propose extensions to the Contract Pattern to include faults. These extensions are elicited using an Audio Visual SoS case study. We discuss how the extended contract pattern relates to our previous work on fault modelling in Section V. Further work and our conclusions are presented in Section VI.

II. Background

Our work builds on the established field of design patterns. Such patterns were originally conceived for designing buildings [8]; their usage has since expanded to software design patterns [9] (particularly for object-oriented software engineering). Recent research provides a number of modelling patterns for SoS architectures [10]. These were described using a structured approach that is underpinned by an "architectural framework" known as the CAFF [11]. The CAFF provides a consistent way of describing both *enabling patterns* ("patterns whose use enables a number of systems engineering applications" [10]) and *architectural frameworks* (which "provide guidance on information sets, or views, that may be used to present architectural information in a standard way" [12]).

This paper builds on two main areas of background work – the Contract Pattern [7] and the Fault Modelling Architectural Framework (FMAF) [13], both of which have been developed using the CAFF. These are described in more depth in Sections II-A and II-B respectively.

Both the Contract Pattern and the FMAF are defined and demonstrated using SysML [14]. They are, however, notation agnostic and several notations have been developed for defining system architectures[1]. We use SysML in our work and in this paper primarily due to its increased use in industry. The profiling features of SysML allow us to extend the base notation with the required concepts for the relevant domain.

A. Contracts for SoS

We have previously defined the Contract Pattern; a modelling pattern for defining contracts in SoSs [7]. The Contract Pattern is considered an enabling pattern in relation to our previous work on modelling patterns [10] – it may be applied in many different domains and across different architectural patterns, for example service-oriented architecture, centralised etc.

The Contract Pattern enables the analysis of emergent behaviours of an SoS, through the modelling of the internal behaviours of CSs. The pattern allows a modeller to define this behaviour in such a way as not to over-constrain the CS implementations. This limited internal behaviour is considered a contract. A CS implementation is expected to conform to one or more contracts.

The Contract Pattern comprises several related viewpoints, defining: (i) the contracts of the SoS; (ii) the conformance of the CSs to the contracts; (iii) the connections between the ports governed by contracts; (iv) the contract definitions; and (v) the protocols constraining the behaviour of the contract. The viewpoints of the Contract Pattern are identified in Table I and are described in more detail in [7]. They assume the use of the Interface Pattern, as defined in [10], when considering the connections between the ports and interfaces governed by contracts.

We propose that contracts defined using the Contract Pattern may be translated to the Compass Modelling Language (CML): a formal modelling language designed for SoSs [15]. Conformance relationships identified between CSs and contracts may be checked in terms of CML trace refinement. Some technical challenges relating to this are discussed in [7].

B. Fault Modelling in SoS

In our previous work we developed a Fault Modelling Architectural Framework (FMAF) for designing fault-tolerant SoSs [16], [13], [17] and demonstrated how fault-tolerant properties of such models could be verified [18][2]. The FMAF defines a set of viewpoints that prompts an SoS developer to consider the impact of faults at the early stages of design, resulting in a coherent set of views that aid the stakeholders of the SoS to understand its erroneous and recovery behaviour.

The FMAF has been developed with respect to established dependability concepts [19]. An SoS *failure* [13] is defined as a deviation of the service provided by the SoS from expected (correct) behaviour. An *error* is defined as the part of the SoS state that can lead to its subsequent service failure. The adjudged or hypothesised cause of an error is called a *fault* and (in keeping with our chosen nomenclature) a failure of a CS can cause a fault of the SoS.

The set of viewpoints prescribed by the FMAF provide an approach for defining: (i) faults, errors and failures of SoSs; (ii) relationships between faults, errors and failures and CSs; (iii) structural designs that enable fault tolerance; (iv) the behaviour of the SoS in the presence of errors; and (v) the recovery behaviour provided by CSs. For further details of these viewpoints, see [17].

III. Related Work

A common theme of both the work on contracts and the work on fault modelling is the value of precision in the definition of interfaces. The SysML language allows basic operation signatures to be defined at interfaces, and pre- and postconditions, though these are rarely used in practice. The Design by Contract (DbC) software engineering technique [20] is built upon to constrain operations in the contract and interface definitions. DbC has been used in several other areas including software engineering design patterns [21] and in service-based components [22]. Another approach to contract design has been proposed in [23]. There contracts specify goals and requirements of an SoS, and are given an LTL semantics. Previous work has considered: nonfunctional properties and DbC in architectural interfaces [24]; how interfaces in SysML can be translated into a formal notation (CML) [15]; and a method of extending SysML interface descriptions with a contractual pattern [6]. The current research builds on this previous work.

Kopetz identifies a variety of potential issues and design principles for the control of cognitive complexity of SoSs [2]. Relied Upon Message Interfaces (RUMIs), similar to the notion of architectural interfaces, are used to specify the exchange of messages between CSs in terms of their "syntax, semantics and temporal behaviour" [2]. Our approach aims to avoid the situation where there are "multiple different semantic specifications of the same RUMI" [2] (one for each CS involved) by providing a standard approach for defining the behaviour of CSs through contract specifications. By providing a notation to define faults and recovery within the contracts pattern we provide support for the identified need to consider "fault containment and error propagation when defining RUMIs" [2].

These are many examples of architecting fault-tolerant systems, however to the best of our knowledge this is limited to the systems level and not in the field of SoS engineering. Relevant work includes: the Error Model Annex for AADL that allows architectural modelling of dependability features [25]; the use of UML for modelling erroneous behaviour in embedded systems [26], [27]; and modelling using SysML for dependable complex physical systems [28] and to verify the safety requirements in the design of embedded safety-critical control systems [29].

IV. Modelling Faults in Contract Descriptions

Bang & Olufsen (B&O) develop home Audio Visual (AV) networks in which several devices cooperate to provide a seamless experience for the user. The AV network exhibits many

[1] In [15] we survey this area more comprehensively, with particular emphasis on the support for the rigorous definition of interface specification.

[2] The FMAF is considered an architectural framework, rather than a pattern, as it is of a larger scale than a pattern, and is designed to cover the full functionality of an SoS.

properties typical of an SoS [1], [5]; for example, the devices (CSs) are *heterogeneous*, with (potentially) a wide variation in *autonomy*. The CSs may include non-B&O products not fully under the control of the SoS. They may be capable of operating independently in a stand-alone mode. Although not distributed widely, the fact that CSs are not co-located offers the possibility of certain *emergent* behaviours, such as sound which "follows" the listener round the home. *Evolution* may be present in the AV SoS as a result of software or firmware upgrades.

In order for the CSs to cooperate effectively they need to establish a single leader to coordinate the SoS. The leader is established through an election process, which involves a series of messages being exchanged between the devices according to a predefined algorithm. For the purposes of this illustrative example we focus on the election functionality of the SoS. To this end, for each *AV Device* we define several contracts including the Leader Election Device (*LE Device*) contract, which specifies the behaviour of the election functionality of an *AV Device*. The composition of contracts making up the SoS is shown in Figure 1.

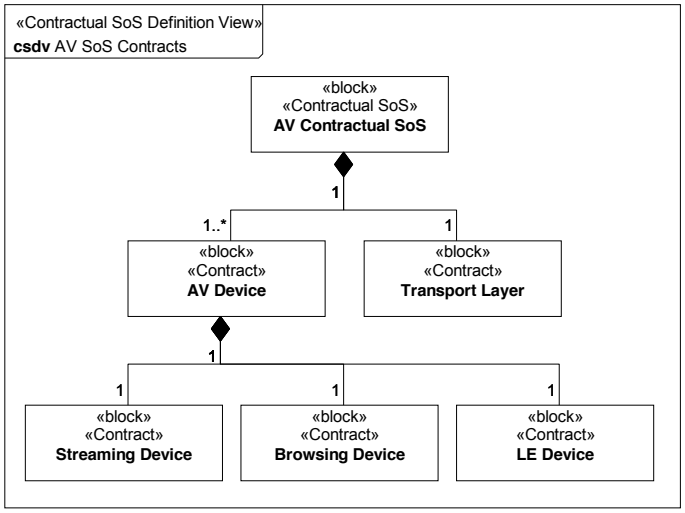

Fig. 1. Contractual SoS Definition View of the AV SoS

The SoS Leader Election protocol [6], [7] also requires a CS for transmitting messages – the behaviour of such a CS is defined in the *Transport Layer (TL)* contract. The reliable transmission of messages is a key aspect of the election. For example, if messages to and from one CS are dropped then that CS may choose to elect a different leader than the rest of the SoS, resulting in an inconsistent AV experience. As the network responsible for transmitting messages is an independent part of the SoS, only limited guarantees of its reliability can be made. Therefore it is important to build fault tolerance into the transmission of messages to tolerate *TL* failures.

Figure 2 shows the state-machine of a well-behaved *TL* (we call the behaviour of a CS/SoS in the absence of faults its *nominal* behaviour). The *TL* maintains a queue of messages in transit. When initialised, an *LE Device* may send a message to the *TL* (*LE_SendMsgs*), in which case the *TL* enters the *Reader* state. The received message is packaged up for the *TL* queue (by the operation *createMessage*), added to the queue, and the *Reader* state is exited. If the queue is not empty the *TL* can enter the *Writer* state, in which a message is taken from the queue (operation *getNext*) and the *TL* attempts to deliver it to its destination. The *Reader* state is exited when either the delivery is successful, or it times out.

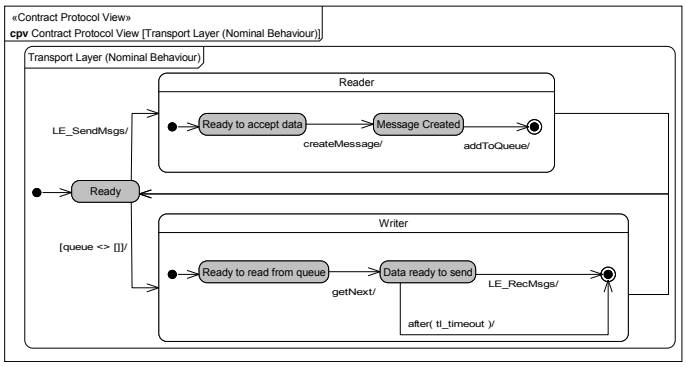

Fig. 2. Contract Protocol View for the Transport Layer

In this case study, we may consider the role of the SoS integrator, and thus observe that the *TL* in fact exhibits erroneous behaviour. The *Faulty TL* is depicted in Figure 3. The *Reader* state is identical, but the *Writer* includes an erroneous transition *dropMessage* (stereotyped as an $\langle\langle error \rangle\rangle$) that drops the message from the *Ready to read from queue* state. This *Faulty TL* roughly corresponds to UDP in that no guarantees can be made over the delivery of messages.

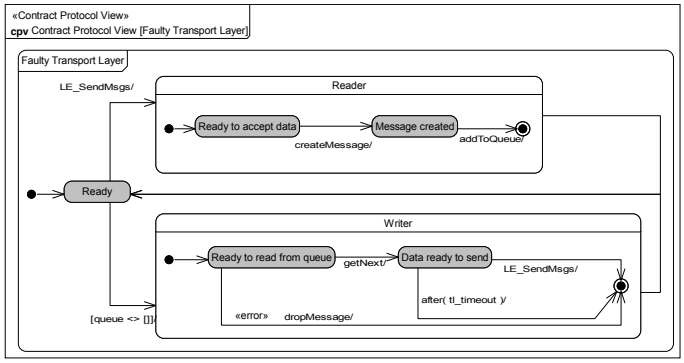

Fig. 3. Contract Protocol View for the Faulty Transport Layer

In this paper we define the behaviour of a wrapper contract that can sit between an *LE Device* and the *TL* in order to tolerate transient[3] faults in the *TL* as depicted in Figure 3. The use of a wrapper contract to ensure the local dependability of a pre-existing contract is a technique advocated by Romanovsky [30].

Figure 4 shows the new *contractual SoS* and the composition of the contracts in the SoS.

Here, there is a single contract describing the behaviour of the *Faulty Transport Layer*, and multiple copies of the *AV*

[3]A fault is *transient* if its "presence is bounded in time" [19]. Note that we cannot tolerate *permanent TL* faults (whose presence are "continuous in time" [19]) as these are indistinguishable from an *LE Device* leaving the SoS (which is part of the nominal behaviour of the SoS).

Device contract, which is in turn made up of three contracts, *Streaming Device*, *Browsing Device* and a *Fault-Tolerant LE Device* contract adapted from the original *LE Device* contract in Figure 1 to compensate for the faults exhibited by the *TL*. The adaption is to compose the basic *LE Device* contract with a set of *LE Wrappers* to create the *Fault-Tolerant LE Device* contract.

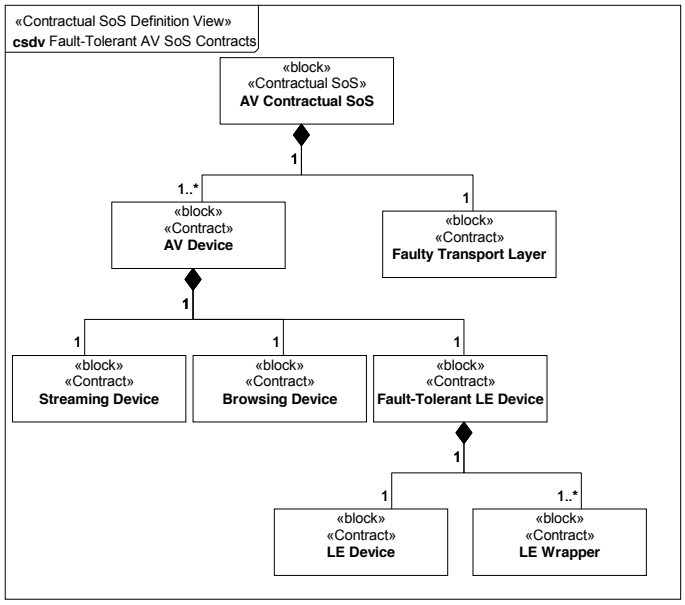

Fig. 4. Contractual SoS Definition View of the Fault-Tolerant AV SoS

The Fault Modelling Contract View in Figure 5 is an instantiation of a new viewpoint, extending the Contract Pattern. The central three blocks are stereotyped as *Fault*, *Error* and *Failure*, nomenclature standardised in [19] and used in the FMAF. Recall, as described in Section II-B that a failure of a CS can cause a fault of the SoS. Each has a textual description. The *Drop Message* error and *Message Loss* failures are artefacts at the CS-level in the *Faulty Transport Layer*. Based upon the dependability concepts used in the FMAF, the CS-level *Message Loss* failure is considered to cause the *Unreliable Message Transmission* fault at the SoS-level. This fault, *located in* the *AV Contractual SoS*, *affects* the *LE Device* contract, and is *mitigated by* the *LE Wrapper* contract.

The purpose of the *LE Wrapper* contract is to acknowledge receipt of messages destined for the *LE Device* contract and to listen for acknowledgments to messages that it has sent. The Contract Definition View in Figure 6 gives the state, operations and invariants of the *LE Wrapper*. Each *Fault-Tolerant LE Device* consists of an *LE Device* and a wrapper for every other *LE Device* with which it communicates. The contract for *LE Wrapper* is therefore parameterised by *myId* (the identity of the "owning" *LE Device*) and *yrId* (the identity of the receiving *LE Device*.) The *snd_pl* value holds the payload that is in the process of being sent at a given time. The figure shows an extension made to the Contract Pattern – the inclusion of a *mitigates* tag, which explicitly records that the *LE Wrapper* contract mitigates the *Unreliable Message Transmission* SoS-level fault, as identified in Figure 5.

The required behaviour of the *LE Wrapper* contract is specified in Figure 7. *Wrapper Send* and *Wrapper Rec* are parallel

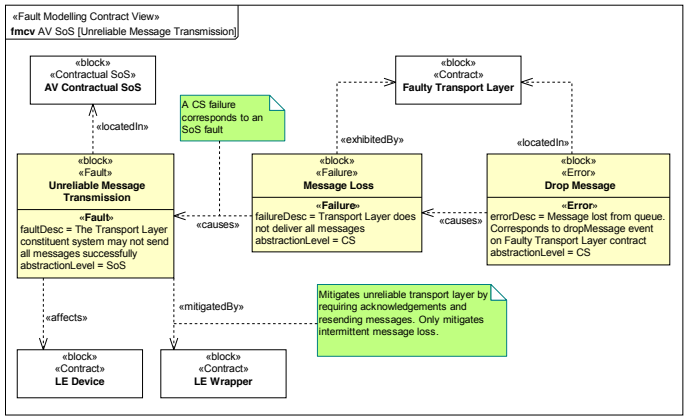

Fig. 5. The Fault Modelling Contract View for the Unreliable Message Transmission fault

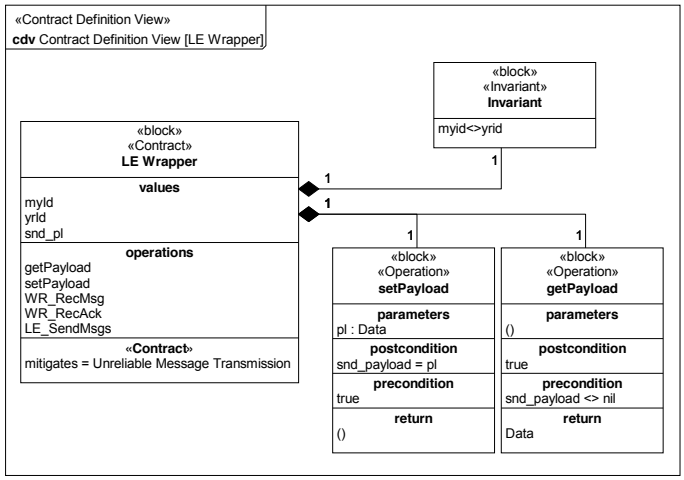

Fig. 6. The Contract Definition View of the LE Wrapper

states. *Wrapper Send* passes on messages from the owning *LE Device* to the *TL*. It begins by storing the message, in case it needs to be resent, then sends it. If an acknowledgment is received (within the time limit *wrapper_timeout*) the wrapper goes back to its initial state, otherwise it tries once again to resend the stored message. This resend only happens once; even if an acknowledgment is not received after the second attempt, the wrapper does not attempt to resend[4]. *Wrapper Rec* receives and acknowledges incoming messages from the *TL* and passes them on to the *LE Device*[5].

In summary, Table I identifies those viewpoints of the Contract Pattern which have been extended to include fault modelling concepts (extensions are shown in italics) and the additional Fault Modelling Contract Viewpoint presented in this paper.

[4]The number of retries can be trivially increased, but the AV SoS needs to be tolerant to devices exiting the SoS at any time, therefore the *LE Wrapper* should not retry to send the message indefinitely. The exact choice on the number of retries is a trade-off between maintaining consistency within the SoS in the presence of faults and reducing unnecessary traffic on the *TL*.

[5]This *LE Wrapper* is therefore similar to, but considerably simpler than, protocols for making UDP reliable such as Automatic Repeat Request (ARQ) and Pragmatic General Multicast (PGM).

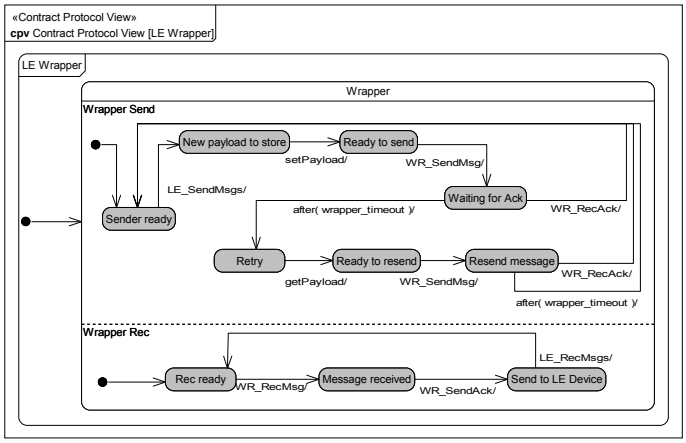

Fig. 7. The Contract Protocol View for the LE Wrapper

TABLE I. INFORMAL DESCRIPTION OF THE EXTENDED CONTRACT PATTERN VIEWPOINTS

Existing Viewpoints	
Name	**Purpose of Viewpoint**
Contractual SoS Definition Viewpoint	Identifies the contracts which comprise the Contractual SoS.
Contract Conformance Viewpoint	Denotes the contracts to which the SoS constituent systems conform.
Contract Connections Viewpoint	Shows connections and interfaces between contracts of the Contractual SoS.
Contract Definition Viewpoint	Defines the operations, state variables and state invariants for a single contract. *Failure modes and mitigated faults added.*
Contract Protocol Viewpoint	Protocol specification of a contract. *Erroneous transitions added and fault-tolerant protocols supported.*
New Viewpoint	
Name	**Purpose of Viewpoint**
Fault Modelling Contract Viewpoint	Associates faults, errors and failures with contracts. Includes *exhibited by*, *mitigated by* and *affected by* relationships.

V. RELATIONSHIP WITH FMAF

As described in Section II, the Contract Pattern and FMAF are defined using a common framework. We may therefore make extensions (as is done with the extended Contract Pattern in this paper) and also combine them using their ontologies to identify shared modelling elements and relationships. In this section, we briefly consider the areas in which the Contract Pattern and FMAF overlap and consider one possible area in which there is a difference in the proposed modelling approach.

The extended Contract Pattern and the FMAF can be used in combination to provide a consistent and complementary set of viewpoints of an SoS model. Both aspects of the model should have a consistent understanding of the faults, errors and failures of the SoS. To achieve this we recommend that a single definition of these is used for both sets of viewpoints, by making use of the FMAF Fault/Error/Failure Definition Viewpoint (FEFDV) [13] within the extended Contract Pattern. The FEFDV provides details of all of the relevant faults, errors and failures of the SoS and identifies any relationships that exist between them. The Fault Modelling Contract View in Figure 5, must therefore be consistent with the FEFDV.

The structural aspects of the patterns are complementary – where the FMAF supports definition of the composition and connections of CSs, the extended Contract Pattern does the same for contracts. The Contract Conformance Viewpoint of the extended Contract Pattern relates the CSs to the contracts and can thus be used as a basis for checking the consistency of these complementary aspects of the model.

The approaches to modelling behaviour in the FMAF and in the extended Contract Pattern differ quite considerably. Whilst the FMAF uses sequence diagrams and activity diagrams to focus on processes of the SoS and interactions between CSs, the extended Contract Pattern uses state machines to target the states and transitions of single CSs (via the definition of the contract protocols they are required to adhere to). These two complementary approaches can be effectively used side by side to provide a more complete treatment of the behaviour of the SoS.

Finally, traceability links could be used as a semi-formal way (see for example [13]) of relating the complementary aspects of the extended Contract Pattern and the FMAF – the implementation of this is the subject of further research.

VI. CONCLUSIONS

In this paper we have highlighted the crucial need to examine the impact of faults when defining contracts for CSs, in order to create a fault-tolerant SoS. Using an Audio Visual SoS case study we have illustrated how the previously developed Contract Pattern [6] can be extended to include consideration for faults and fault tolerance. Extending the Contract Pattern in this way has been shown to be complementary to our previous efforts in defining an architectural framework (the FMAF [13]) for incorporating faults at the architectural design phase of creating an SoS.

Fault modelling has been an active area of research for some time and there is a rich taxonomy of fault types, as well as a rich set of "countermeasures" (fault tolerance and recovery, mitigation, etc) [19]. In this paper we have used a single fault and recovery mechanism in order to investigate the feasibility of integration of the Contract Pattern and the FMAF. A natural continuation of this work is to extend the types of faults and countermeasures that we cover, as well as considering the interaction of multiple faults. Our preliminary results may also be strengthened by their application to a wider range of SoSs, leading to a more complete set of extensions to the Contract Pattern for the purpose of modelling faults and fault tolerance within SoSs.

The integration of the extended Contract Pattern with the FMAF also deserves a more thorough treatment, and we consider two main directions: the first of which relates to the first piece of future work mentioned earlier – identifying new viewpoints that may be needed to support a wider variety of faults and countermeasures, along with new relationships between such viewpoints and the FMAF; and the second is the recording of traceability links (as discussed in Section V) between elements of the extended Contract Pattern and elements of the FMAF (building upon the work in [13]).

Further types of extension to the Contract Pattern could be envisaged. Ongoing work on static fault analysis of an SoS investigates marking up structural architectural models with certain fault information. Analysis can then be carried

out using an industrial fault analysis tool. It seems natural to incorporate the mark-up language within the extended Contract Pattern.

ACKNOWLEDGEMENTS

The work presented here is supported by the EU Framework 7 Integrated Project "Comprehensive Modelling for Advanced Systems of Systems" (COMPASS, Grant Agreement 287829) (for more information see http://www.compass-research.eu) and the EPSRC Platform Grant on Trustworthy Ambient Systems (TrAmS-2).

REFERENCES

[1] M. W. Maier, "Architecting Principles for Systems-of-Systems," *Systems Engineering*, vol. 1, no. 4, pp. 267–284, 1998.

[2] H. Kopetz, "System-of-systems complexity," in *Proc. 1st Workshop on Advances in Systems of Systems*, K. G. Larsen, A. Legay, and U. Nyman, Eds., 2013.

[3] Y. Mordecai and D. Dori, "A model-based framework for architecting system-of-systems interoperability, interconnectivity, interfacing, integration, and interaction," in *23rd INCOSE International Symposium.*, 2013.

[4] B. W. Clifton, S. Hostetler, and W. N. Felder, "Mathematical models of emergence in complex systems-of-systems," in *23rd INCOSE International Symposium.*, 2013.

[5] C. B. Nielsen, P. G. Larsen, J. Fitzgerald, J. Woodcock, and J. Peleska, "Model-based engineering of systems of systems," *In preparation for submission to ACM Computing Surveys*, January 2013.

[6] J. Bryans, J. S. Fitzgerald, R. Payne, and K. Kristensen, "Maintaining Emergence in Systems of Systems Integration: a Contractual Approach using SysML," in *Annual International Symposium of the International Council on Systems Engineering*. INCOSE, June-July 2014, to appear.

[7] J. Bryans, J. Fitzgerald, R. Payne, A. Miyazawa, and K. Kristensen, "SysML Contracts for Systems of Systems," June 2014, IEEE Systems of Systems Engineering Conference 2014.

[8] C. Alexander, M. Ishikawa, M. Silverstein, M. Jacobson, I. Fiksdahl-Ling, and S. Angel, *A Pattern Language*. Oxford University Press, 1977.

[9] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, *Design Patterns Elements of Reusable Object Oriented Software*. Addison Wesley, 1995.

[10] S. Perry, "Report on Modelling Patterns for SoS Architectures," COMPASS Deliverable, D22.3, Tech. Rep., February 2013. [Online]. Available: http://www.compass-research.eu/deliverables.html

[11] J. Holt, S. Perry, F. O. Hansen, S. Hallerstede, and K. Kristensen, "Initial Report on Guidelines for Architectural Level SoS Modelling," COMPASS Deliverable, D21.2, Tech. Rep., 2013, available at http://www.compass-research.eu/.

[12] J. Holt, C. B. Nielsen, F. O. Hansen, A. Miyazawa, R. Payne, J. Iyoda, and M. Cornelio, "Initial Report on SoS Architectural Models," COMPASS Deliverable, D22.1, Tech. Rep., March 2012, available at http://www.compass-research.eu/.

[13] Z. Andrews, C. Ingram, R. Payne, A. Romanovsky, S. Perry, and J. Holt, "Traceable Engineering of Fault-Tolerant SoSs," in *International INCOSE Symposium, Las Vegas, USA,*, June-July 2014, to appear.

[14] OMG, "OMG Systems Modelling Language Version 1.3," Available: http://www.omg.org/spec/SysML/1.3 (Accessed June 2013), June 2012.

[15] J. Bryans, R. Payne, J. Holt, and S. Perry, "Semi-formal and formal interface specification for system of systems architecture," in *7th International Systems Conference, IEEE SysCon*. IEEE, April 2013.

[16] Z. Andrews, J. Fitzgerald, R. Payne, and A. Romanovsky, "Fault Modelling for Systems of Systems," in *Proceedings of the 11th International Symposium on Autonomous Decentralised Systems (ISADS 2013)*, March 2013, pp. 59–66.

[17] Z. Andrews, A. Didier, R. Payne, C. Ingram, J. Holt, S. Perry, M. Oliveira, J. Woodcock, A. Mota, and A. Romanovsky, "Report on timed fault tree analysis – fault modelling," COMPASS Deliverable, D24.2, Tech. Rep., September 2013. [Online]. Available: http://www.compass-research.eu/deliverables.html

[18] Z. Andrews, R. Payne, A. Romanovsky, A. L. Didier, and A. Mota, "Model-based development of fault tolerant systems of systems," in *2013 IEEE International Systems Conference (SysCon 2013)*, Orlando, USA, April 2013.

[19] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, "Basic Concepts and Taxonomy of Dependable and Secure Computing," *IEEE Transactions on Dependable and Secure Computing*, vol. 1, pp. 11–33, 2004.

[20] B. Meyer, *Object-Oriented Software Construction (2nd ed.)*. Prentice-Hall, 1988.

[21] M. D. Champlain, "The contract pattern," in *In Proceedings of Pattern Languages of Program Design 4 (PLoPD4*, 1997.

[22] A. Beugnard, J.-M. Jézéquel, and N. Plouzeau, "Making components contract aware," *IEEE Computer*, vol. 32, no. 7, pp. 38–45, 1999.

[23] A. Arnold, B. Boyer, and A. Legay, "Contracts and Behavioral Patterns for SoS: The EU IP DANSE approach," in *AiSoS*, ser. EPTCS, K. G. Larsen, A. Legay, and U. Nyman, Eds., vol. 133, 2013, pp. 47–66.

[24] R. J. Payne and J. S. Fitzgerald, "Evaluation of Architectural Frameworks Supporting Contract-based Specification," School of Computing Science, Newcastle University, Tech. Rep. CS-TR-1233, December 2010.

[25] A.-E. Rugina, P. H. Feiler, K. Kanoun, and M. Kaâniche, "Software dependability modeling using an industry-standard architecture description language," *CoRR*, vol. abs/0809.4109, 2008.

[26] S. Bernardi and J. Merseguer, "A UML profile for dependability analysis of real-time embedded systems," in *Proceedings of the 6th international workshop on Software and performance*, ser. WOSP '07. New York, NY, USA: ACM, 2007, pp. 115–124.

[27] A. Bondavalli, M. D. Cin, D. Latella, I. Majzik, A. Pataricza, and G. Savoia, "Dependability analysis in the early phases of UML-based system design." *Comput. Syst. Sci. Eng.*, vol. 16, no. 5, pp. 265–275, 2001.

[28] P. David, V. Idasiak, and F. Kratz, "Reliability study of complex physical systems using SysML," *Reliability Engineering System Safety*, vol. 95, no. 4, pp. 431–450, 2010. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0951832009002671

[29] J.-F. Pétin, D. Evrot, G. Morel, and P. Lamy, "Combining SysML and formal models for safety requirements verification," in *ICSSEA 2010 — 22nd International Conference on Software & Systems Engineering and their Applications*, 2010.

[30] A. B. Romanovsky, "Exception handling in component-based system development," in *25th International Computer Software and Applications Conference*, 2001, pp. 580–586.